



10/14/99

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
REQUEST FOR FILING NATIONAL PATENT APPLICATION
under 35 USC 111(a) and Rule 53(b) WITH SIGNED DECLARATION

1
PATENT
APPLICATION

Asst. Commissioner for Patents
BOX PATENT APPLICATION
Washington, D.C. 20231

NONPROVISIONAL

Sir:

Enclosed herewith is the PATENT APPLICATION of

Inventor: Angshuman SAHA, et al.

Title: Method of Link Word Synchronization

(Our Deposit Account No. 03-3975)

Our Order No.	71575	0239604
Client #		Matter #
Atty. Docket	239604	PL-012
TMC#		Client Ref

1c542 U.S. PTO
09/418119
10/14/99

including:

Date : October 14, 1999

1. ☒ Specification, claims and abstract: **25** pages
2. ☒ Declaration ☒ Original ☐ Facsimile/Copy
3. ☒ Drawings: **2** sheet(s) ☒ informal ☐ formal
4. ☒ An Assignment and cover sheet. Please return the recorded Assignment to the undersigned.
5. ☒ A Verified Statement establishing "small entity" status under Rules 9 & 27.

THE FOLLOWING FILING FEE IS BASED ON CLAIMS AS FILED LESS ANY ABOVE CANCELLED

6. Basic Filing Fee				\$760/\$380	\$	380.00	101/201
7. Total Claims:	41	minus 20 =	21	x \$18/\$9 =	+	189.00	103/203
8. Independent Claims:	06	minus 3 =	03	x \$78/\$39 =	+	117.00	102/202
9.	TOTAL FILING FEE ENCLOSED =				\$	686.00	
10. If "assignment" box is X'd, add recording fee				+ \$40/\$40	+	40.00	581
TOTAL FEE ENCLOSED:						\$	726.00

CHARGE STATEMENT: The Commissioner is hereby authorized to charge any missing or insufficient fees which may be required relative to this application, or credit any overpayment, to Deposit Account 03-3975, for which purpose a duplicate copy of this sheet is enclosed.

1100 New York Avenue, N.W.
Ninth Floor, East Tower
Washington, D.C. 20005-3918
Tel: (650) 233-4790
Fax: (650) 233-4545
DAJ/g

PILLSBURY MADISON & SUTRO LLP
By: David A. Jakopin, Reg. No. 32,995

Express Mail Label:
Date of Deposit:

EK 025 838 705 US
October 14, 1999

I certify that this paper and listed enclosures are being deposited with the U.S. Post Office
"Express Mail Post Office to Addressee" under 35 CFR 1.10 on the above date, addressed
to Asst. Commissioner for Patents, Box Patent Application, Washington, D.C. 20231

Valerie J. Harmon

Inventors: Angshuman SAHA, et al.
App. No.: Unassigned
Filed: Herewith

Atty Dkt. 71575/239604
Client Ref: PL-012

Title: **Method of Link Word Synchronization**

VERIFIED STATEMENT (DECLARATION) CLAIMING SMALL ENTITY
STATUS (37 CFR 1.9(d) and 1.27(c)) - SMALL BUSINESS CONCERN

I hereby declare that I am an official empowered to act on behalf of the small business concern identified below:

NAME OF CONCERN: **PLURIS, INC.**
ADDRESS OF CONCERN **10455a Bandley Drive, Cupertino, CA 95014**

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.12, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees under Section 41(a) and (b) of Title 35, United States Code, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention entitled as above and invented by **SAHA**, et al. described in the above-captioned specification.

If the rights held by the above-identified small business concern are not exclusive, each small entity, individual, concern or organization having rights to the invention is listed below* and no rights to the invention are held by any person, other than the inventor, who could not qualify under 37 CFR 1.9(c) as an independent inventor if that person had made the invention, or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d) or a nonprofit organization under 37 CFR 1.9(e). *Note: Separate verified statements are required from each person, concern or organization having rights to the invention, averring to small entity status (37 CFR 1.27).

FULL NAME of _____

☐ INDIVIDUAL ☒ SMALL BUSINESS CONCERN ☐ NONPROFIT ORGANIZATION

I acknowledge the duty to file, in this case, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF SIGNATORY: Bulent Erbilgin
TITLE: Vice President, Engineering
ADDRESS: 10455A Bandley Drive
Cupertino, CA 95014

SIGNATURE



DATE

9/9/1999

Method of Link Word Synchronization

BACKGROUND OF THE INVENTION

1. Field of Invention

The present invention relates to a method for synchronization for use in data transmission and more particularly to a method for word synchronization that can be used in systems with multiple serial links.

2. Description of Related Art

The problem of word synchronization results in communications systems with multiple serial links where each link carries a portion of a word being transmitted. For example, a 64-bit word may be transmitted between two word devices as eight 8-bit bytes, the arrival of which must be synchronized so that the 64-bit word arrives correctly. In order to enable this synchronization in a robust way, requirements for handshaking go beyond what has been developed for a single serial link. For example, the Next Generation I/O Link Architecture Specification describes link establishment and maintenance handshaking together with an associated state machine but only for a single serial link. There is no generalization to the concept of word synchronization on multiple serial links. (NGIO Specification, Oct. 30, 1998, Intel Corporation; http://developer.intel.com/design/servers/future_server_io/ for an overview, and http://developer.intel.com/design/servers/future_server_io/documents/ngiolas.pdf for the complete document.)

Likewise, Fibre Channel specifications provide link states and link establishment handshaking but only for a single serial link. (ANSI X3.230:1994, Fibre Channel Physical and Signaling Interface (FC-PH) Rev 4.3)

Only limited word synchronization has been developed for multiple serial devices. The Vitesse quad SERDES (serializer/deserializer) device provides the capability to synchronize multiple serial links but without any status, confirmation or link state handshake, features that are essential to a robust implementation. (Vitesse Quad SERDES VSC7216 specification; Vitesse Quad SERDES VSC7214 Specification: <http://www.vitesse.com/summarysheets/vsc7214.htm>.)

For example, in the absence of handshaking, data may be sent without synchronization and errors may develop and go undetected.

There are many ways of encoding data to send it over a serial data link. The most widely used is an 8-bit to 10-bit encoding defined in the Fibre Channel PH Specification cited above.

This encoding allows the receiver to recover a clock for sampling the data and to detect byte boundaries in the data. It also provides for the transmission of non-data control characters. While Fibre Channel defines a use for some control characters, they can be assigned other uses when the encoding is used separately from Fibre Channel. There are other less widely used encodings with similar features. These encodings, including that used by Fibre Channel and others, will be referred to hereinafter as FC-like serial encoding. A serializer or deserializer that uses an FC-like serial encoding will hereinafter be called a SERDES.

SUMMARY OF THE INVENTION

Accordingly, it is an object of this invention to provide a method of word synchronization in a system of word devices connected by multiple data links.

It is a further object of this invention to provide a method of word synchronization in a system of word devices connected by multiple data links that are otherwise independent and unsynchronized.

It is a further object of this invention to provide word synchronization in a system of word devices including SERDES devices connected by serial lines.

It is a further object of this invention to provide robust word synchronization in a system of word devices so that a loss of synchronization leads to a re-establishment of synchronization.

5 It is a further object of this invention to provide positive confirmation of word synchronization with each packet of data transmitted between two word devices.

It is a further object of this invention to minimize the number of overhead bytes transmitted between two word devices in addition to the data being transmitted.

10 The above and related objects are realized according to a first aspect of the present invention by a method for word synchronization between a first word device and a second word device. Initially the first word device and the second word device are each unsynchronized. The first word device and the second word device each include serializers and deserializers. The serializers of the first word device are connected to the deserializers of the second word device by multiple serial lines, and the serializers of the second word device are connected to the
15 deserializers of the first word device by multiple serial lines. The method includes:

transmitting a synchronization request on the serial lines from the serializers of the first device to the deserializers of the second device;

becoming synchronized at the second device in response to the synchronization request from the first device;

20 transmitting a synchronization request on the serial lines from the serializers of the second device to the deserializers of the first device after the second device has become synchronized;

becoming synchronized at the first device in response to the synchronization request from the second device; and

transmitting data on the serial lines from the serializers of the first device to the deserializers of the second device after the first device has become synchronized in response to the request for synchronization from the synchronized second device.

This aspect of the invention solves the problem of word synchronization between two communicating devices that are using multiple high-speed serial lines with clock information embedded in the data. Other approaches have been restricted to single serial links or to limited word synchronization. The present invention can be applied to systems with multiple high-speed serial lines. It will work in rugged environments where there is loss of power and cable failure. The present invention allows recovery from loss of synchronization due, for example, to transmission errors and leads to a return to word synchronization.

In a preferred embodiment the method also includes detecting loss of synchronization between the first and second devices and confirming synchronization with the transmission and reception of each packet of data. This allows the method to respond quickly to the loss of synchronization and to maintain an accurate status of the link.

These and other objects and advantages of the invention will become more apparent and more readily appreciated from the following detailed description of the presently preferred exemplary embodiment of the invention taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram illustrating a SERDES system according to the present invention; and

FIG. 2 is state transition diagram according to the present invention.

**DETAILED DESCRIPTION OF THE
PRESENTLY PREFERRED EXEMPLARY EMBODIMENT**

A preferred embodiment of the present invention provides a method of achieving word
5 synchronization between two asynchronous devices connected by multiple serial links. As
illustrated in FIG. 1, word devices 2 and 4 transfer data to one another on a 64-bit bi-directional
link bus 3 preferably running at 125 Mhz. The data is actually transmitted via a set of eight
SERDES (serializer/deserializer) devices, each of which converts 8 bits of parallel data to a bit-
serial form in which data and clock are encoded in a single serial signal that is transmitted across
10 a serial line. Correspondingly the data is received via a set of eight SERDES devices, each of
which converts data in bit serial form to 8 bits of parallel data. Thus, in FIG. 1 word device 2 is
connected to eight serializers 6a – 6h for transmitting data and eight deserializers 8a – 8h for
receiving data. Likewise, word device 4 is connected to eight serializers 10a – 10h for
transmitting data and eight deserializers 12a – 12h for receiving data. Transmission from device
15 2 to device 4 occurs across serial lines 14a – 14h, and transmission from device 4 to device 2
occurs across serial lines 16a – 16h.

Collectively, serial lines 14a – 14h and 16a – 16h are denoted as the fabric link (or simply
the link) between word devices 2 and 4. In the preferred embodiment serial lines 14a – 14h and
16a – 16h are optical links running at 1.25 Gb/sec; alternatively, the transmission may be
20 electronic rather than optical. Additionally, ribbon fiber may be used to carry all of the serial
lines in a single connector. Although FIG. 1 illustrates a case where the link consists of eight
serial links (i.e., serial lines 14a – 14h and 16a – 16h), the method can be applied generally to a
link consisting of N serial links. Although the method is described using serial links, it can be
applied to any set of links that deliver data with different delays.

A single 64-bit data unit transmitted from device 2 or device 4 constitutes a word, that is, a 64-bit word. Each word must be correctly delivered to the other device in a single 64-bit unit. Accomplishing this requires synchronizing the arrival of the eight 8-bit bytes. When this synchronization is in effect, the link (including serial links 14a – 14h and 16a – 16h) is word
 5 synchronized.

SERDES devices 6a – 6h, 8a – 8h, 10a – 10h, and 12a – 12h are consistent with the specifications of Vitesse Quad SERDES devices (e.g., VSC7214/VSC7215). These SERDES devices are able to transmit a special sequence of Idle characters, synchronized and in parallel on multiple serial links. Upon receiving such a sequence on multiple serial links, these SERDES
 10 devices are capable of synchronizing the serial lines so that bytes simultaneously input to the transmitting SERDES devices (e.g., 6a – 6h) are then simultaneously output from the receiving SERDES devices (e.g., 12a – 12h). This feature makes it possible to achieve word synchronization. However, SERDES devices 6a – 6h, 8a – 8h, 10a – 10h, and 12a – 12h do not include a mechanism for indicating whether a link between two word devices is word synchronized.

15 Furthermore, when such a link is not word synchronized, there is no mechanism for requesting from one end of the link to the other that the special synchronization-enabling Idle sequence be sent. Each of SERDES devices 6a – 6h, 8a – 8h, 10a – 10h, and 12a – 12h does have error signals for data errors and loss of byte synchronization on the receiving end of a serial link.

Although the presence of these error signals does indicate the loss of word synchronization, their
 20 absence does not imply word synchronization. FIG. 2 illustrates the ability of receiving SERDES devices 8a – 8h and 12a – 12h to send a loss of synchronization (LOS) signal to word units 2 and 4 respectfully.

A word link between two word devices must establish word synchronization in both directions before it can operate normally. If word synchronization is lost, it must be reestablished before normal operations can resume. Word synchronization requires that all the serial links that comprise the fabric link are working in unison to correctly deliver each 64-bit word intact (i.e., all bytes delivered in the same clock cycle). In order to accomplish word synchronization between devices 2 and 4, several operations must be carried out in a coordinated way as each word device operates as a transitter or as a receiver for a 64-bit word. Acting as a transmitter, a device must send a data stream that enables the corresponding receiver to establish word synchronization. Acting as a receiver, a device must detect when word synchronization has been established or lost, notify a corresponding transmitter when there is no word synchronization, and request from a corresponding transmitter that a sequence to enable word synchronization be sent.

Synchronization is achieved between devices 2 and 4 according to the state transition diagram of FIG. 2. The possible link states for devices 2 and 4 are described in Table 1: no_synch, half_synch, synch, and active. As indicated in FIG. 2 in parentheses for each state, the device output sequence associated with each state is transmitted across the link to the device thereby connected (e.g., from device 2 to device 4). The possible device output sequences (i.e., sequences of words) are described in Table 2. Wsynch_A denotes a sequence of sixteen idles followed by WSA, a defined control character. Wsynch_B denotes a sequence of sixteen idles followed by WSB, a defined control character. In addition to these output sequences specifically used to coordinate word synchronization, output can include properly formatted packets and idles.

Control words for the system illustrated in FIG. 1 and FIG. 2 are defined in terms of the control characters specified for a single SERDES device. An enumeration of defined control

characters is given in Table 3. WSA and WSB are control characters defined by the SERDES device to achieve byte synchronization. SERDES packet sequencing is facilitated by control characters for start of packet (SOP), end of packet (EOP), plus end and start of packet (ESOP). SERDES devices also use Idle characters to maintain byte synchronization and Abort characters to abort a packet. For any of these control characters, the corresponding control word is that control character repeated across each serial line. In FIG. 2, a bad control word denotes a properly formatted control word where the bytes are not identical across each serial line. Additionally a control word to denote loss of synch (LOS) effectively results from an LOS signal from any SERDES as illustrated in FIG. 1.

As illustrated in FIG. 2, state transitions are associated with receipt of a control word (or an error in a control word). State transition arrows in FIG. 2 are labeled with control word inputs that cause that state change; inputs that do not cause a state change are not shown.

When a word device (e.g., device 2 or 4) is in the no_synch state, this indicates that the receiver at this end of the link does not have word synch. That is, the device is not prepared to synchronously receive 64-bit words. This is the default state when the device is first powered up. The output generated, which is a request for word synchronization, is Wsynch_A, a 16-word sequence of idles followed by the control word WSA. This 17-word output sequence is sent repeatedly as long as the device remains in the no_synch state. When the device receives control word WSA it switches to the half_synch state. When the device receives control word WSB, it sends the control sequence Wsynch_B and then switches to the synch state. The no_synch state is reached whenever the device detects loss of synch (LOS).

When a word device is in the half_synch state, this indicates that the device at this end of the link has word synch so that the device at the other end of the link may begin sending packets.

However, there is no confirmation that the device at the other end of the link has word synch.

The output generated is Wsynch_B, a 16-word sequence of idles followed by the control word

WSB. This 17-word output sequence is sent repeatedly as long as the device remains in the

half_synch state. It enables the device at the other end of the link to become synchronized, and

5 notifies that device that this device has achieved word synchronization. When the device

receives control words WSB or SOP, it switches to the synch state. When the device receives

control word LOS or a bad control word, it switches to the no_synch state. The half_synch state

is reached whenever the device receives WSA.

When a word device is in the synch state, devices at both ends of the link have word

10 synch, and packets can be sent in either direction. In this state the device sends available packets

of data across the link. However, the link is not in confirmed good state either because no packet

has been received at this end yet or because a bad control word has been received. (Receipt of a

bad control word indicates a possible loss of word synch.) When the device receives control

word SOP, it switches to the active state. When the device receives control word WSA, it

15 switches to the half_synch state. When the device receives control word LOS or a bad control

word, it switches to the no_synch state.

When a word device is in the active state, the link is in a confirmed state of word synch in

the sense that both ends have word synch and the last SOP received was good. In this state the

device sends available packets of data across the link. When the device receives a bad control

20 word, it switches to the synch state. When the device receives control word LOS or a bad control

word, it switches to the no_synch state.

For example, assume word devices 2 and 4 are reset simultaneously so that each device is

put into the no_synch state. Further assume that devices 2 and 4 are capable of normal (i.e.,

error-free) operation after this reset. Then each device outputs Wsynch_A. Each device then receives control word WA and is put into the half_synch state so that each device outputs Wsynch_B. Each device then receives control word WB and is put into the synch state so that each device begins sending normal data, which includes the control word SOP at the start of a packet. Each device then receives control word SOP and is put into active state and continues sending normal data.

Next assume that just device 2 is reset so that device 2 is put into the no_synch state while device 4 is in the active state. Device 2 outputs Wsynch_A while device 4 is sending normal data. Device 4 then receives WSA and is put into the half_synch state so that it outputs Wsynch_B. When device 2 receives WSB, it is put into the half_synch so that it outputs Wsynch_B. Then, as discussed above, devices 2 and 4 can both proceed from the half_synch state to the synch state and then to the active state.

Similarly, if devices 2 and 4 are in the active state and the transmission is disrupted by a loss of synch (LOS) or a bad control word, then devices 2 and 4 can both return to the active state with confirmed word synchronization and normal transmission of data. The present invention robustly enables word synchronization by confirming the synchronization of the transmission after a disruption and the subsequent resumption of normal operations.

Preferably the method, as illustrated by the state transition diagram of FIG. 2, is implemented as an ASIC (Application-Specific Integrated Circuit) that maintains the states of word devices 2 and 4 and controls their output. However, the method may also be implemented by other means including a general-purpose computing device although such an implementation may be undesirably slow.

Preferably, the fabric link does include a CPU (Central Processing Unit) that can be used for monitoring the status of the link. This CPU may receive a maskable interrupt whenever the link enters the no_synch state, and possibly also the half_synch state. These features enable the local CPU to monitor the state of the link in order to take appropriate action when the status of the link changes, thereby diagnosing and working around link failures and directing data traffic toward or away from the link as appropriate. The capability to monitor the status of the link and to diagnose and resolve failures is especially desirable when the method is applied in complex networks with multiple word devices.

When a word device is in the synch state or the active state, additional aspects of the method enable the confirmation of synchronization with the transmission and reception of each packet of data so that a loss of synchronization manifests itself by the receipt of one or more bad control words. For example, according to the details of a network implementation, a designated control word (e.g., SOP) may be preceded and followed by some specified number of words including idles, data, and other control words, but not the same designated control word. By not allowing any two copies of this control word to be sent without some minimal number of other intervening words, it can be ensured that the receiver only receives that control word when word synchronized. If the receiver were not synchronized, it could mix bytes of the designated control word with other nearby words, but no such mixing could produce a copy of the designated control word because of the separation required in the sending of these words. Therefore, reception of the designated control word is positive confirmation of link synchronization. It is possible to have more than one such control word.

More generally, the use of control characters only as elements of properly formed control words provides detection of loss of synchronization even in the absence of a required minimal

separation between repeated control words. That is, a properly formed control word includes the same control character on each serial line, where that control character is only used in this way.

If any of these characters is ever received without being part of a corresponding control word, then this receipt of a bad control word is either the result of a loss of synchronization or a data
 5 error.

Data errors are generally transient, and so their repeated detection over a short time interval indicates a probable loss of synchronization. For example, according to the preferred embodiment shown in FIG. 2, when a word device in the active state receives a first bad control word, the device transitions to the synch state. From the synch state, receipt of a second bad
 10 control word results in a transition to the no_synch state; however, receipt of an SOP control word without errors returns the device to the active state. Thus, according to the preferred embodiment, a loss of synchronization results from the receipt of two bad control words without an intervening SOP control word, which indicates the synchronized transmission of data. This principle can be generalized, for example, by counting the number of received bad control words
 15 up to some threshold before transitioning from the synch state to the no_synch state. In this way, the invention allows for a faster recovery from transient data errors at the expense of a reduced sensitivity to a loss of synchronization.

Using control words to confirm synchronization according to the present invention is consistent with their functional use in data transmission. In the preferred embodiment, suitable
 20 control words defined by identical control characters across each serial line are used as framing indicators, including (from Table 3): SOP (start of packet), EOP (end of packet), and ESOP (end and start of packet). These allow confirmation of link synchronization at the start and end of

every packet received. The control characters used in these control words are not used for any other purpose, so these control words also enable the detection of a loss of synchronization.

Although only a single exemplary embodiment of this invention has been described in detail above, those skilled in the art will readily appreciate that many modifications are possible in the exemplary embodiment without materially departing from the novel teachings and advantages of this invention. Accordingly, all such modifications are intended to be included within the scope of this invention.

Table 1 Link **State Definitions**

<u>State</u>	<u>Description</u>
no_synch	<ul style="list-style-type: none"> Receiver at this end of link does not have word synch. This is the default state on power up and reset. Device sends Wsynch_A sequence to request word synchronization. Device switches to half_synch state when it receives control word WSA. Device sends Wsynch_B and switches to synch state when it receives control word WSB. Device goes to this state whenever it detects loss of word synch (LOS).
half_synch	<ul style="list-style-type: none"> Receiver at this end of link has word synch. Device at the other end of link may begin sending packets. Device sends Wsynch_B sequence to request word synchronization. Device switches to synch state when it receives control words WSB or SOP. Device switches to no_synch state when it receives control word LOS or a bad control word. Device goes to this state whenever it receives WSA.
synch	<ul style="list-style-type: none"> Receivers at both ends of the link have word synch Packets can be sent in either direction. However, the link is not in confirmed good state either because no packet has been received at this end yet or because a bad control word has been received. (Receipt of a bad control word indicates a possible loss of word synch.) Device switches to active state when it receives control word SOP. Device switches to half_synch state when it receives control word WSA. Device switches to no_synch state when it receives control word LOS or a bad control word.
active	<ul style="list-style-type: none"> Link is in a confirmed state of word synch: both ends have word synch and last SOP received was good. Device switches to synch state when it receives a bad control word. Device switches to no_synch state when it receives control word LOS or a bad control word.

Table 2 Output Sequences

<u>Output</u>	Description
Wsych_A	Special sequence of sixteen idles followed by WSA.
Wsych_B	Special sequence of sixteen idles followed by WSB.
data (normal op)	Properly framed packets and idles.

Table 3 SERDES Control Character Definitions

<u>Control Char. Name</u>	Comment	Can legally follow these:	Can legally precede these:
WSA – Word Sych A	WSA is sent to request a word sych sequence when word sych is lost	Word sych enabling Idle sequence.	Word sych enabling Idle sequence.
WSB – Word Sych B	WSB is sent to indicate the ability to accept packets and to facilitate attaining word sych when a receiver has word sych but the receiver at the other end might not.	Word sych enabling Idle sequence.	Idle, SOP.
SOP – Start of Packet	SOP identifies the start of a packet.	WSB, Idle, EOP.	Packet, starting with fabric header.
EOP – End of Packet	EOP identifies the end of packet	Packet, ending with CRC.	Idle, SOP.
ESOP – End and Start of Packet	ESOP separates back-to-back packets thereby replacing EOP-SOP control word sequence.	Packet.	Packet.
Idle	Idles are periodically required by SERDES to maintain byte and word sych. Idles can be sent before SOP, after EOP, and after WSB.	Word sych enabling Idle sequence, WSA.	Word sych enabling Idle sequence, WSA, WSB.
Abort	Abort is sent to terminate the transmission of a packet. When sent within a packet (i.e., between SOP and EOP, or equivalently with ESOP), Abort causes the packet it interrupts to be dropped (or delivered to a CPU that has requested back packets). When sent between packets, it has no effect. Implementations may (but are not required to) consider Abort to be a bad control word, so more than one Abort without an intervening SOP may cause the link to revert to no_sych state.	Anything except WSA.	Idle, SOP.

What is claimed is:

1. A method for word synchronization between a first word device and a second word device, the first word device and the second word device each being unsynchronized, the first word device and the second word device each including a plurality of serializers and a plurality of deserializers, the serializers of the first word device being connected to the deserializers of the second word device by a plurality of serial lines, and the serializers of the second word device being connected to the deserializers of the first word device by a plurality of serial lines, the method comprising:

transmitting a synchronization request on the plurality of serial lines from the serializers of the first device to the deserializers of the second device;

becoming synchronized at the second device in response to the synchronization request from the first device;

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device;

becoming synchronized at the first device in response to the synchronization request from the second device; and

transmitting data on the plurality of serial lines from the serializers of the first device to the deserializers of the second device after the first device has become synchronized in response to the request for synchronization from the synchronized second device.

2. The method of claim 1, further comprising: transmitting data on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the synchronized second device has received data transmitted from the first device.

3. The method of claim 1, further comprising: becoming unsynchronized at the first device in response to receiving a loss of synch signal from at least one of the deserializers of the first device.

4. The method of claim 1, further comprising:

detecting at the second device a bad control word that is inconsistent across the deserializers of the second device; and

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the second device has detected the bad control word.

5. The method of claim 1, wherein the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

6. The method of claim 5, further comprising:

detecting at the second device a bad control word that is inconsistent across the deserializers of the second device; and

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the second device has detected the bad control word.

7. The method of claim 1, wherein the serial lines are fiber optics.

8. The method of claim 1, wherein a ribbon fiber includes the serial lines.

9. A method for word synchronization between a first word device and a second word device, the first word device being unsynchronized, the second word device being synchronized, the first word device and the second word device each including a plurality of serializers and a plurality of deserializers, the serializers of the first word device being connected to the deserializers of the second word device by a plurality of serial lines, and the serializers of the second word device being connected to the deserializers of the first word device by a plurality of serial lines, the method comprising:

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device;

becoming synchronized at the first device in response to the synchronization request from the second device;

transmitting data on the plurality of serial lines from the serializers of the first device to the deserializers of the second device after the first device has become synchronized in response to a request for synchronization from the synchronized second device.

10. The method of claim 9, further comprising: transmitting data on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the synchronized second device has received data transmitted from the first device.

11. The method of claim 9, further comprising: becoming unsynchronized at the first device in response to receiving a loss of synch signal from at least one of the deserializers of the first device.

12. The method of claim 9, further comprising:

detecting at the second device a bad control word that is inconsistent across the deserializers of the second device; and

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the second device has detected the bad control word.

13. The method of claim 9, wherein the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

14. The method of claim 13, further comprising:

detecting at the second device a bad control word that is inconsistent across the deserializers of the second device; and

transmitting a synchronization request on the plurality of serial lines from the serializers of the second device to the deserializers of the first device after the second device has detected the bad control word.

15. The method of claim 9, wherein the serial lines are fiber optics.

16. The method of claim 9, wherein a ribbon fiber includes the serial lines.

17. A method for word synchronization between a plurality of word devices connected by a plurality of serial lines, comprising the steps of:

requesting synchronization from a first device to a second device when the first device does not have synchronization;

requesting synchronization from a first device to a second device when the first device has synchronization;

receiving a request for synchronization at a first device from a second device, the first device then becoming synchronized;

transmitting data from a first device to a second device, the first device being synchronized, the first device having received from the second device a synchronization signal indicating that the second device is synchronized.

18. The method of claim 17, wherein, in the step of transmitting data, the synchronization signal includes at least one of a synchronization request from the synchronized second device and a start-of-packet indicator from data transmitted by the second device.

19. The method of claim 17, further comprising: becoming unsynchronized at a device in response to receiving a loss-of-synch signal.

20. The method of claim 19, wherein the loss-of-synch signal is generated by a deserializer included in the device.

21. The method of claim 17, further comprising:
 detecting a bad control word at a first device from a second device; and
 requesting synchronization from a first device to a second device, the first device having received a bad control word from the second device.
22. The method of claim 17, wherein
 the word units include serializers and deserializers that satisfy a SERDES specification for control characters,
 a bad control word received by a device is inconsistent across deserializers of the device.
23. A method for detecting and adapting to a loss of word synchronization at a first word device, the first word device being synchronized and connected to a second word device by a plurality of serial lines, the method comprising: becoming unsynchronized at the first device in response to receiving a loss of synch signal from at least one of the serial lines connected to the second device.
24. The method of claim 23, wherein
 the first word device and the second word device each include a plurality of serializers and deserializers;
 the serial lines connect the serializers of the first word device to the deserializers of the second word device and the serializers of the second word device to the deserializers of the first word device; and

the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

25. The method of claim 23, further comprising: becoming unsynchronized at the first device in response to receiving a threshold number of bad control words from the serial lines connected to the second device, wherein consecutive bad control words are not separated by a synchronized data packet.

26. The method of claim 25, wherein the threshold number of bad control words is one.

27. The method of claim 25, wherein the threshold number of bad control words is greater than one.

28. The method of claim 25, wherein
the first word device and the second word device each include a plurality of serializers and deserializers;

the serial lines connect the serializers of the first word device to the deserializers of the second word device and the serializers of the second word device to the deserializers of the first word device; and

the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

29. The method of claim 28, wherein the threshold number of bad control words is one.

30. The method of claim 28, wherein the threshold number of bad control words is greater than one.

31. A method for detecting and adapting to a loss of word synchronization at a first word device, the first word device being synchronized and connected to a second word device by a plurality of serial lines, the method comprising: becoming unsynchronized at the first device in response to receiving a threshold number of bad control words from the serial lines connected to the second device, wherein consecutive bad control words are not separated by a synchronized data packet.

32. The method of claim 31, wherein the threshold number of bad control words is one.

33. The method of claim 31, wherein the threshold number of bad control words is greater than one.

34. The method of claim 31, wherein

the first word device and the second word device each include a plurality of serializers and deserializers;

the serial lines connect the serializers of the first word device to the deserializers of the second word device and the serializers of the second word device to the deserializers of the first word device; and

the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

35. The method of claim 34, wherein the threshold number of bad control words is one.

36. The method of claim 34, wherein the threshold number of bad control words is greater than one.

37. A method for confirming word synchronization at a first word device, the first word device being actively synchronized and connected to a second word device by a plurality of serial lines, the method comprising: detecting a confirming control word from the serial lines connected to the second device, wherein the second device separates the transmission of confirming control words by a threshold number of intervening words.

38. The method of claim 37, wherein the threshold number of intervening words is sufficient to prevent an unsynchronized reception of the designated control word at the first word device.

39. The method of claim 37, wherein
the first word device and the second word device each include a plurality of serializers and deserializers;

the serial lines connect the serializers of the first word device to the deserializers of the second word device and the serializers of the second word device to the deserializers of the first word device; and

the serializers and the deserializers of the first and second devices satisfy a SERDES specification for control characters.

40. The method of claim 39, wherein the threshold number of intervening words is sufficient to prevent an unsynchronized reception of the designated control word at the first word device.

41. The method of claim 39 wherein the confirming control word is also used as a framing indicator.

ABSTRACT OF THE DISCLOSURE

A method for word synchronization can be applied to asynchronous devices including SERDES devices connected across serial lines. A state transition methodology characterizes the state of the device based on control characters received consistently across the serial lines and channels the system to a state of word synchronization. Loss of synchronization and transmission errors lead to a re-establishment of synchronization.

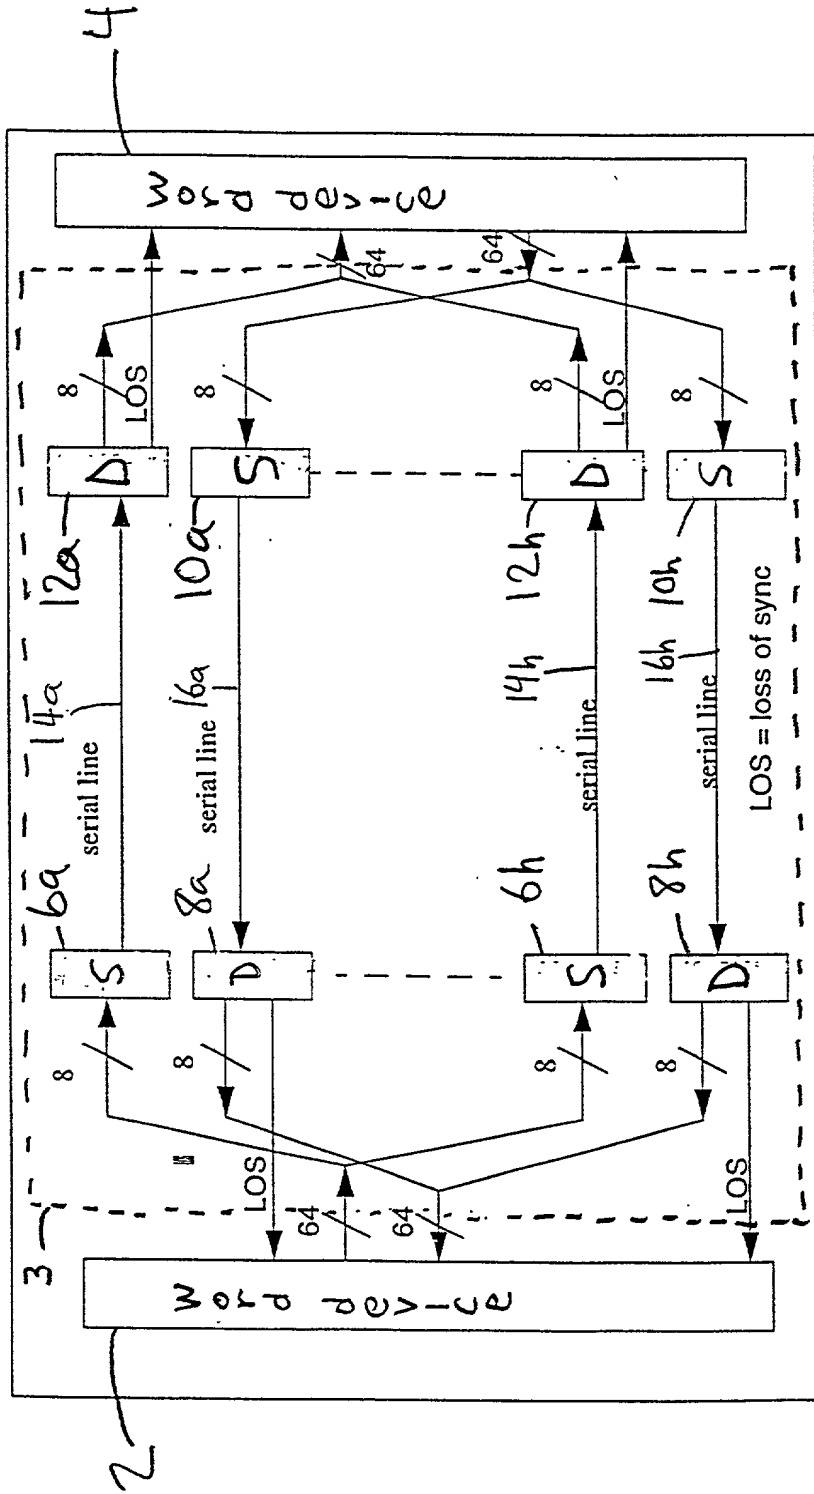
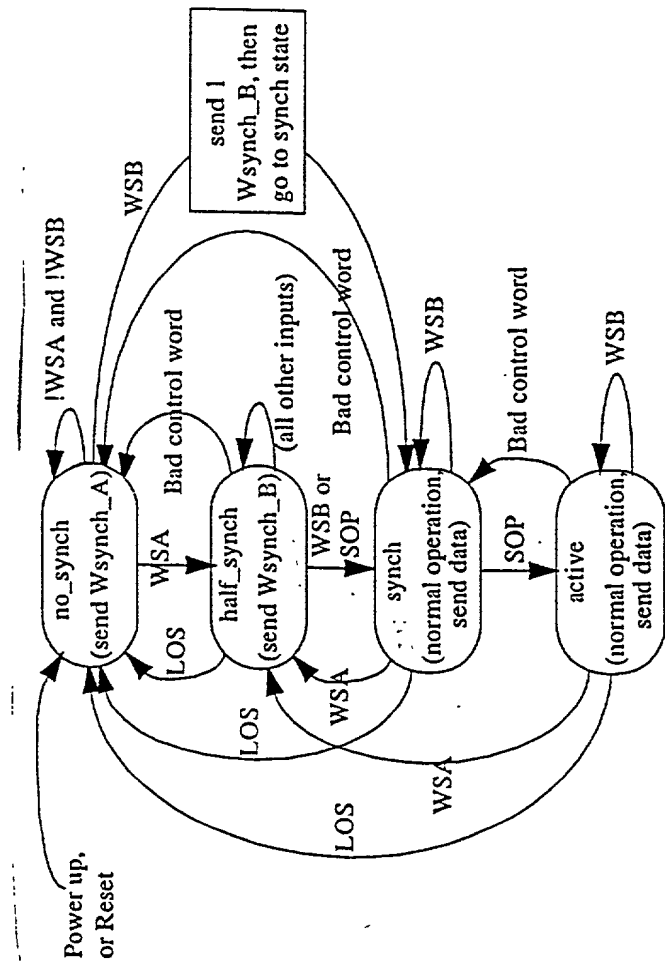


FIG. 1



X! : Anything but X

FIG. 2

RULE 63 (37 C.F.R. 1.63)
DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

As a below named inventor, I hereby declare that my residence, post office address and citizenship are as stated below next to my name, and I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled, **Method of Link Word Synchronization**, the specification of which [X] is attached hereto, bearing Atty Docket No. **71575/239604**

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose all information known to me to be material to patentability as defined in 37 C.F.R. 1.56. I hereby claim foreign priority benefits under 35 U.S.C. 119/365 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate filed by me or my assignee disclosing the subject matter claimed in this application and having a filing date (1) before that of the application on which priority is claimed, or (2) if no priority claimed, before the filing date of this application:

PRIOR FOREIGN APPLICATION(S):		Date first Laid-	Date Patented	Priority Claimed
<u>Number</u>	<u>Country</u>	<u>open or Published</u>	<u>or Granted:</u>	<u>Yes</u> <u>No</u>

I hereby claim domestic priority benefit under 35 U.S.C. 119/120/365 of the indicated United States applications listed below and PCT international applications listed above or below and, if this is a continuation-in-part (CIP) application, insofar as the subject matter disclosed and claimed in this application is in addition to that disclosed in such prior applications, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in 37 C.F.R. 1.56 which became available between the filing date of each such prior application and the national or PCT international filing date of this application:

PRIOR U.S. PROVISIONAL, NONPROVISIONAL AND/OR PCT APPLICATION(S)		Status	Priority Claimed?
<u>Application No.:</u>	<u>Day/MONTH/Year Filed:</u>	<u>(pending, abandoned, patented)</u>	<u>Yes</u> <u>No</u>

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

And I hereby appoint Pillsbury Madison & Sutro LLP, 1100 New York Avenue, N.W., Ninth Floor, East Tower, Washington, D.C. 20005-3918, telephone number (650) 233-4790 (to whom all communications are to be directed), and the below-named persons (of the same address) individually and collectively my attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith and with the resulting patent, and I hereby authorize them to delete persons no longer with their firm and to act and rely on instructions from and communicate directly with the person/assignee who first sent this case to them and by whom I hereby declare that I have consented after full disclosure to be represented unless/until I instruct the above Firm and/or a below attorney in writing to the contrary.

Paul N. Kokulis	16773	Peter W. Gowdey	25872	David A. Jakopin	32995	Roger R. Wise	31204
Raymond F. Lippitt	17519	Dale S. Lazar	28872	Mark G. Paulson	30793	Jay M. Finkelstem	21082
G. Lloyd Knight	17698	Glenn J. Perry	28458	Timothy J. Klima	34852	Anita M. Kirkpatrick	32617
Carl G. Love	18781	Kendrew H. Colton	30368	Stephen C. Glazier	31361	Michael R. Dzwonczyk	36787
Kevin E. Joyce	20508	Paul E. White, Jr.	32011	Paul F. McQuade	31542	W. Patrick Bengtsson	32456
George M. Sirilla	18221	G. Paul Edgell	24238	Ruth N. Morduch	31044	Jack S. Barufka	37087
Donald J. Bird	25323	Lynn E. Eccleston	35861	Richard H. Zaitlen	27248	Adam R. Hess	41835

1. INVENTOR'S SIGNATURE: *Angshuman SAHA*
Inventor's Name: **Angshuman SAHA**
Residence (City, State): **Sunnyvale, California**
Post Office Address: **471 Acalanes Drive, #1, Sunnyvale, CA 94086**

Date: 10/11/99
Country of Citizenship: **India**

2. INVENTOR'S SIGNATURE: *Steven Farnworth*
Inventor's Name: **Steven FARNWORTH**
Residence (City, State): **Menlo Park, California**
Post Office Address: **764 Live Oak Avenue, Menlo Park, CA 95138**

Date: 10/12/99
Country of Citizenship: **United States of America**

3. INVENTOR'S SIGNATURE: *Russell R. Tuck, III*
Inventor's Name: **Russell R. TUCK, III**
Residence (City, State): **San Jose, California**
Post Office Address: **1136 Blaney Avenue, San Jose, CA 95129**

Date: 10/11/99
Country of Citizenship: **United States of America**

4. INVENTOR'S SIGNATURE: *Deepak Mansharani*
Inventor's Name: **Deepak Mansharani**
Residence (City, State): **San Jose, California**
Post Office Address: **1738 Commodore Drive, San Jose, CA 95133**

Date: 10/11/99
Country of Citizenship: **United States of America**